

Catalogue no. 11-622-M — No. 025

ISSN: 1705-6896

ISBN: 978-1-100-20914-2

Research Paper

The Canadian Economy in Transition Series

Firm Dynamics: Employment Growth Rates of Small Versus Large Firms in Canada

by Jay Dixon and Anne-Marie Rollin

Economic Analysis Division

18th Floor, R.H. Coats Building, 100 Tunney's Pasture Driveway,
Ottawa, Ontario K1A 0T6

Telephone: 1-800-263-1136



Statistics
Canada

Statistique
Canada

Canada

How to obtain more information

Specific inquiries about this product and related statistics or services should be directed to the Media Hotline, Communications and Library Services Division, Statistics Canada, Ottawa, Ontario K1A 0T6 (telephone: 613-951-4636).

For information about this product or the wide range of services and data available from Statistics Canada, visit our website at www.statcan.gc.ca or contact us by e-mail at infostats@statcan.gc.ca or by telephone from 8:30 a.m. to 4:30 p.m. Monday to Friday:

Statistics Canada National Contact Centre

Toll-free telephone (Canada and the United States):

Inquiries line	1-800-263-1136
National telecommunications device for the hearing impaired	1-800-363-7629
Fax line	1-877-287-4369

Local or international calls:

Inquiries line	1-613-951-8116
Fax line	1-613-951-0581

Depository services program

Inquiries line	1-800-635-7943
Fax line	1-800-565-7757

Information to access the product

This product, Catalogue no. 11-622-M, is available for free in electronic format. To obtain a single issue, visit our website at www.statcan.gc.ca under "Our agency" click on Site map > Statistics and studies > and select "Publications".

Standards of service to the public

Statistics Canada is committed to serving its clients in a prompt, reliable and courteous manner. To this end, the Agency has developed standards of service which its employees observe in serving its clients. To obtain a copy of these service standards, please contact Statistics Canada toll free at 1-800-263-1136. The service standards are also published on www.statcan.gc.ca. Under "Our agency" click on About us > The agency > and select "Providing services to Canadians".

The Canadian Economy in Transition

The Canadian Economy in Transition is a series of new analytical reports that investigate the dynamics of industrial change in the Canadian economy. This new series brings together a coherent set of research reports that provide users with a wide variety of empirical perspectives on the economy's changing industrial structure. These perspectives include the dynamics of productivity, profitability, employment, output, investment, occupational structure and industrial geography. Readers are encouraged to contact the authors with comments, criticisms and suggestions.

All papers in **The Canadian Economy in Transition** series go through institutional and peer review to ensure that they conform to Statistics Canada's mandate as a government statistical agency and adhere to generally accepted standards of good professional practice.

The papers in the series often include results derived from multivariate analysis or other statistical techniques. It should be recognized that the results of these analyses are subject to uncertainty in the reported estimates.

The level of uncertainty will depend on several factors: the nature of the functional form used in the multivariate analysis; the type of econometric technique employed; the appropriateness of the statistical assumptions embedded in the model or technique; the comprehensiveness of the variables included in the analysis; and the accuracy of the data that are utilized. The peer group review process is meant to ensure that the papers in the series have followed accepted standards to minimize problems in each of these areas.

Statistics Canada
Economic Analysis Division

Firm Dynamics: Employment Growth Rates of Small Versus Large Firms in Canada

By Jay Dixon and Anne-Marie Rollin

Published by authority of the Minister responsible for Statistics Canada

© Minister of Industry, 2012

All rights reserved. Use of this publication is governed by the Statistics Canada Open Licence Agreement (<http://www.statcan.gc.ca/reference/copyright-droit-auteur-eng.htm>).

July 2012

Catalogue no. 11-622-M, no. 025
Frequency: Occasional

ISSN 1705-6896
ISBN 978-1-100-20914-2

Ottawa

Authors' names are listed alphabetically.

La version française de cette publication est disponible (n° 11-622-M au catalogue, n° 025).

Note of appreciation

Canada owes the success of its statistical system to a long-standing partnership between Statistics Canada, the citizens of Canada, its businesses, governments and other institutions. Accurate and timely statistical information could not be produced without their continued cooperation and goodwill.



Acknowledgements

The authors would like to thank John Baldwin for his support in this research project, Annette Ryan and Larry Shute of Industry Canada, and Leonard Landry for help with the LEAP database. Javier Miranda and Robert Petrunia provided helpful comments.


Symbols

The following standard symbols are used in Statistics Canada publications:

- . not available for any reference period
- ... not available for a specific reference period
- ... not applicable
- 0 true zero or a value rounded to zero
- 0^s value rounded to 0 (zero) where there is a meaningful distinction between true zero and the value that was rounded
- ^p preliminary
- ^r revised
- x suppressed to meet the confidentiality requirements of the *Statistics Act*
- E use with caution
- F too unreliable to be published
- * significantly different from reference category ($p < 0.05$)

Table of contents

Abstract	6
Executive summary	7
1 Introduction.....	8
2 Conceptual issues and literature review	9
2.1 Size and growth	9
2.2 The HJM critique	11
2.2.1 Data issues: Firms versus establishments	12
2.2.2 Regression-to-the-mean	12
2.2.3 Firm age effect.....	13
2.3 Firm size distributions.....	14
3 The LEAP database	15
3.1 Measures of employment	16
3.2 Defining age.....	18
3.3 Empirical framework: Growth	19
3.4 Empirical framework: Size classification.....	19
3.5 Econometric framework.....	20
4 Results.....	21
4.1 Descriptive statistics.....	21
4.1.1 Distribution of firms and employment across size and age classes	21
4.1.2 Mean growth rate by size and age	25
4.2 Regression results – Conditional means	28
4.3 Size distribution.....	38
5 Conclusion	40
6 Appendix	42
6.1 Comparison between Average Labour Unit (ALU) and Individual Labour Unit (ILU)	43
References	45



Abstract

This paper examines whether Canadian firms of different sizes (in terms of employment) grow at different rates year-on-year. The data are from Statistics Canada's Longitudinal Employment Analysis Program and cover the 1999-to-2008 period. The methodology is similar to that used by Haltiwanger, Jarmin and Miranda (2010) for the United States: controls are used for firm age, and possible bias from short-term regression to the mean is removed by sizing firms according to their average number of employees in both previous and current years.

The analysis shows that employment growth rates across the Canadian business sector does not vary much between firms of different size classes, except for the smallest and youngest firms. Employment growth rates rise with firm size for firms with fewer than 20 employees, but for larger firms, no relationship emerges between employment growth and firm size. These results are consistent with the average proportionate growth condition of Gibrat's Law—the assertion of French economist Robert Gibrat that average employment growth is independent of firm size.

More studies related to [small-firms Analysis](#) and to [labour market dynamics](#) are available in [Update on Economical analysis](#) (www.statcan.gc.ca/economicanalysis).



Executive summary

Who creates jobs? The popular perception is that it is small firms. Empirical evidence on this issue can be found in studies that examine whether employment growth rates are independent of firm size, a result known as the average proportionate growth condition of Gibrat's Law—the assertion of French economist Robert Gibrat, in 1931, that average firm growth is independent of size.

While some studies have found at least a weak negative relationship between firm size and job growth, in a recent working paper, Haltiwanger, Jarmin and Miranda (2010) maintain that the negative relationship between firm size and job growth reported in some earlier studies is actually a product of inappropriate data, omitted variables, and statistical fallacies. Using the U.S. Census Bureau's Longitudinal Business Database, Haltiwanger et al. correct for these problems and find no systematic relationship between firm size and annual employment growth.

This paper uses Statistics Canada's Longitudinal Employment Analysis Program (LEAP) database to determine if the result is the same for Canada over the 1999-to-2008 period. A major problem with previous studies is that jobs that change hands through mergers, acquisitions and divestures could not be distinguished from the creation of new positions and the destruction of old ones. Also, most datasets do not contain accurate information on the age of firms, which Haltiwanger et al. argue is an important determinant of employment growth. However, it is possible to remove merger and acquisition activity from the LEAP data, at least on a year-over-year basis, and measure firm age in a consistent way.

The methodology in this study is similar to that used by Haltiwanger et al.: controls are included for firm age, and possible bias from short-term regression to the mean is removed by sizing firms according to their average number of employees in the previous and current year. With these adjustments, average annual employment growth rates rise (rather than decrease) with firm size for firms with less than 20 employees. And for firms with 20 employees or more, no relationship is evident between employment growth and firm size. This latter pattern is consistent with the average proportionate growth condition of Gibrat's Law.

These findings are extended by looking at the distribution of firm sizes in the economy. The distribution of firms by size accords with Zipf's Law, which implies that small firms are common whereas large firms are rare, another indication of the applicability of Gibrat's Law to the business sector in Canada.



1 Introduction

Who creates jobs? While the popular perception is that the answer is small firms,¹ the literature is inconclusive. Some studies find a negative relationship between firm size and net job creation, but others do not.

The question is important since job creation is often a political and social imperative. Especially in recessions, but also when unemployment is deemed too high, decision-makers focus on strategies to increase employment. Knowing where jobs are being created is valuable for policy formulation.

As well, the employment growth patterns of firms provide general information about the economy. Typically, a large number of small firms coexist with a small number of large firms (Bartelsman et al., 2003). Studying these distributions and the processes that generate them can inform models of the functioning of the economy.²

Some studies have found at least a weak negative relationship between firm size and job growth, which would seem to confirm the popular wisdom. However, others argue in favour of the average proportionate growth condition of Gibrat's Law—the assertion, in 1931, of French economist Robert Gibrat that employment growth rates are independent of firm size.

The question of which firms create jobs can be approached in two ways. The first is to examine how the rate of employment growth is related to firm size. The second is to determine if some groupings of firms, by virtue of their weight in the economy, are responsible for more net job creation. Small firms can experience higher growth, yet be too small a fraction of the economy to be important job creators. Large firms may grow slowly, but because of their dominance, still be where most of the jobs are to be found. If the rate of employment growth is relatively constant across size groupings, then the share of aggregate job growth attributable to each size group is proportional to its importance in the economy as an employer.

This paper focuses on the first approach. The relative importance of small and large firms in the Canadian economy is the subject of future research.

Using Statistics Canada's Longitudinal Employment Analysis Program (LEAP) database, this paper explores whether annual rates of employment growth of different-sized firms in the Canadian business sector are more or less the same. The analysis is conducted at level of the business sector as a whole over a 10-year period. Two different measures of employment are used: one that corrects for job quality and one that does not.

Section 2 reviews the literature on the importance and implications of the size-growth relationship. Section 3 describes the structure of the LEAP database and its analytical advantages and disadvantages. Section 4 presents the findings of the analysis: descriptive statistics and regression results. Section 5 contains the conclusions of the study and suggests directions for future research.

1. Firm size is measured by the number of employees. Other possibilities include value of assets, sales, etc., but as Sutton (1997) notes, these dimensions of size are rarely explored, at least in the size/growth literature.

2. For an overview of theories consistent with growth, see Geroski (2005) and Sutton (1997).



2 Conceptual issues and literature review

Birch (1981) is often credited with popularizing the importance of small firms for job creation. His claim that small firms (20 employees or less) were responsible for two-thirds of all net jobs created in the United States in the 1970s captured the imagination of small business advocates and politicians and set off a vigorous debate about the validity of his results.

Birch's findings contradicted the growth condition stated in Gibrat's "law of proportional effect": that observed employment growth rates and firm size are independent.³ Gibrat's Law is in many ways a (disputed) regularity in search of a theory. Its main implication is that whether they are large or small, firms face similar idiosyncratic shocks, the incidence and magnitude of which are independent of firm size. Firm growth rates are a consequence of these shocks. A corollary is that a firm's capacity to exploit or deal with whatever these shocks bring is positively related to its size. If it holds, Gibrat's Law should manifest itself in the relationship between firm size and employment growth rates, and in the shape of the distribution of firm size.

Accordingly, studies of the independence of growth rates fall into two categories. The more extensive literature looks directly at the relationship between firm size and employment growth rates. The other looks at the firm-size distribution, and infers whether the growth process that generates it suggests a relationship between firm size and employment growth. Both approaches are employed in this report.

Like most empirical studies, this analysis focuses on mean growth rates across firm size classes in order to address the question regarding which firms create proportionally more jobs. But the full description of the stochastic process associated with Gibrat's Law that is used to infer long-run distributions also assumes a constant variance of the distribution of outcomes (Mata, 2008). The similarity of shocks and/or outcomes across firm sizes that is reflected in average growth rates does not necessarily apply to higher moments of their distributions. And while the question of whether growth is more volatile at the small or large end of the firm size distribution is important, it is beyond the scope of the present paper.

2.1 Size and growth

The literature on firm size and employment growth is old and extensive, but researchers disagree on what it shows. Geroski (2005) claims that fifty-plus years of firm growth research suggest randomness.⁴ On the other hand, in their summary of a half-century of empirical work, Audretsch et al. (2004) mostly document departures from Gibrat's Law for a range of countries over different periods.

3. Mathematically, if S_{t-1} is a firm's size in period $t-1$, then $S_t = (1 + E_{t-1}) * S_{t-1}$, where E is an independent and identically distributed shock.

4. At least, that is how they appear to researchers, if not to the firms themselves. An interpretation is that given the information to which researchers, and presumably policy-makers, have access, it is impossible to identify which firms are destined to grow at any given time.

This literature contains three testable propositions (Petrunia, 2008): that mean employment growth rates are independent of firm size; that the variance of employment growth rates is independent of firm size; and that shocks are uncorrelated over time.⁵ However, the set of firms to which these propositions of size-independent processes and size-dependent reactions should apply is not clear.

It is sometimes assumed that the propositions apply only to continuing firms, with firm births and deaths treated separately. That is, conditional on existence, growth opportunities are independent of firm size, and the probability that not-yet existing firms would take advantage of them is governed by another process.⁶

Alternatively, the propositions could apply to all firms. Ideally, the growth of dying firms would be measured as the rates they would have posted had they survived. The rates for entering firms, on the other hand, could be calculated according to some "if all goes well" optimal or forecast size. Of course, measuring would-be's and would-have-been's presents obvious difficulties. Surveys have been used, but the precision of the estimates derived from these sources is unknown. And this information is not available in administrative databases such as the LEAP. Without such information, entry rates will always be unbounded, because initial size is zero, or if created using an arbitrary sizing, likely to be unnaturally large. Rates of decline associated with deaths will be censored and subject to measurement error.

Another approach is to assume that Gibrat's Law applies only to firms—entering, continuing or exiting—that have reached minimum efficient scale (MES). Employment growth rates are hypothesized to be independent of firm size, except for small, young firms; if they were not born at MES, they will either grow rapidly to it, or fail.⁷ Under this interpretation, selection bias may cause the data to show a relationship between firm size and employment growth even if employment growth is independent of greater-than-MES size, provided the probability of failure is not (Audretsch et al., 2004). If small firms with negative or slow growth rates are more likely to fall below MES and thus exit, the sample of small firms will be biased towards surviving high growth firms while the data will record the full range of growth rates for larger firms.

The literature that has addressed the "who creates jobs" question covers a range of countries over many time periods. In addition to the United States, work has been done for Canada (Baldwin and Picot, 1994; Baldwin, 1998), the United Kingdom, Italy, Portugal, Germany, Japan, and Greece. Some studies contain data that go back to the 19th century, while others exploit contemporary sources. These studies generally examine only the mean employment growth/firm size relationship. Many of them either reject independence in favour of a negative relationship between firm size and employment growth, or find the evidence to be mixed across detailed industry classifications. Some studies include results for variances, showing that they are not independent of firm size, but rarely delve deeper.

Until recently, research has been constrained by limitations of the data that were available. Early studies rely on databases of publicly traded companies, in which small firms are under-represented. Most use only firm-level data, although some cover plants (Baldwin and Picot, 1994; Baldwin, 1998). Plant studies allow for examination of the growth characteristics of the

5. Petrunia (2008) finds that most Canadian industries violate at least one of these three conditions. Although he uses a related dataset (T2-LEAP), his study differs from this one in important respects. For example, his results are for detailed industries, whereas the results of this analysis average over the entire economy. Second, he measures growth over a nine-year period and includes growth due to mergers, whereas year-on-year growth is used here and excludes growth due to mergers. We remove mergers because they do not "create" jobs—at least not in the short run. Petrunia finds that many industries have similar mean growth rates across size groups, but almost all industries have variances that differ across size classes and exhibit an autocorrelation of growth.

6. The probability that a new firm takes advantage of an opportunity is often set arbitrarily (Sutton, 1997).

7. And since the probability that entrants arrive at sizes less than MES varies across industries (Baldwin and Gu, 2011), the applicability of Gibrat's Law also varies across industries.

basic building blocks of firms; firm studies have problems separating out the characteristics of growth at this most basic level from growth that is due to mergers and acquisitions.

In addition, most studies deal exclusively with the manufacturing sector (Baldwin and Haltiwanger, 1998). Audretsch et al. (2004) argue that findings contradicting Gibrat may apply to manufacturing, where entrants are often smaller than MES. Many manufacturing industries entail large sunk costs, require substantial capital investments or benefit from economies of scale. These industries are more favourable to the survival of large firms, and are particularly perilous for slower-growing firms that are below the MES boundary.

By contrast, industries that are easier to enter, for example, in the service sector, are more likely to treat large and small laggards equally.⁸ Audretsch et al. (2004) find that many of the sub-industries of the Dutch service sector exhibit fewer deviations from the law of proportionate growth, as predicted.⁹ Petrunia (2008), on the other hand, finds that in Canada, retail and manufacturing firms both fail to comply with all the requirements of Gibrat's Law—in particular, the constancy of variance across size classes. Neumark et al. (2008) find a similar symmetry in growth dynamics between the manufacturing and service sectors.

The results of individual studies may be affected by their coverage of industry and time. Theory and evidence on expanding industries suggest that small firms grow initially, but fail as the industry matures, leaving mostly larger firms. Declining industries, on the other hand, will likely display different dynamics. For instance, large firms may shrink as the industry shrivels, and be the first to fail when they become small. Thus, the relationship between firm size and employment growth that prevails on average may differ from that for individual industries at different periods in their evolution as they face changing macroeconomic conditions.

As well, there is little evidence suggesting that the properties and effects of technology and other shocks confronting firms are the same across industries and time. In some industries, small firms may exploit changes more easily than do less nimble large firms. Also, a relationship between firm size and growth in a particular industry during a specific period may not persist.

This paper ignores the differences across sectors and poses the question at the level of the business sector as a whole over a relatively long time period, recognizing that macroeconomic cyclical events may well impact differentially on firms of different sizes and that industry-specific technological shocks may also drive growth rates temporarily away from Gibrat's general tendency. It asks whether at a higher level of generality—across industries and over time—the rates of growth of different sized firms in the Canadian business sector are more or less the same for the time period examined.

2.2 The HJM critique

A 1993 study by Davis, Haltiwanger and Schuh and a 2010 working paper by Haltiwanger, Jarmin and Miranda (henceforth referred to as HJM) question the negative relationship between firm size and employment growth found in much research. For three reasons, they contend that the results may be misleading. First, the datasets used are problematic. The data usually fail to distinguish between employment growth resulting from organic internal job creation and job growth attributable to changes in industry structure. Second, the data exhibit regression-to-the-mean effects that should be taken into account. And third, they argue that size effects may be confused with the effects of firm age.

8. See Baldwin and Gu (2011) and Baldwin and Lafrance (2011) for comparisons of the entry process across goods and service industries.

9. Luttmer (2007) develops a model explicitly linking growth to entry costs and to the ease of imitating incumbents. He documents the conditions under which the resulting distribution approaches Zipf's Law.

HJM use the U.S. Census Bureau's Longitudinal Business Database to address these problems. The authors maintain that when using better data, and when controlling for regression-to-the-mean and age, the negative relationship between firm size and employment growth disappears.

2.2.1 Data issues: Firms versus establishments

Davis et al. (1993) and HJM observe that most datasets pertain either to firms or to establishments, and therefore, pertain to different components of the firm and therefore may present different pictures of the process of growth. According to HJM, neither is sufficient on its own to address the size/growth issue.

For some purposes, firms are often more interesting than establishments when it comes to growth since they are the economic decision-making unit. HJM use the example of the retail industry, where growth occurs through the opening of new establishments, rather than the expansion of existing ones. Establishment-level studies, if used alone, show job growth in these industries comes mostly from births, and job destruction from deaths; firm studies provide an extra dimension because overall firm growth involves both birth and death, which must be jointly evaluated to capture the relationship between employment growth and firm size.

However, the firm-level data that are often used make some studies problematic. Employment changes at the establishment level are well defined—all jobs that are added or subtracted are "organic" job creation or job destruction. By contrast, job changes at the firm level may include not only "organic" changes, but also false births, deaths, expansions and contractions associated with mergers, acquisitions and divestitures. These actions, many of which may amount to little more than a shuffling of existing jobs, may create the appearance of job growth when none has actually occurred. In some databases, a firm selling a plant to another firm may appear as the contraction of one and expansion of the other, and a firm spinning-off one of its parts may appear as the simultaneous contraction of a large firm, and the birth of a smaller one.

This merger-and-acquisition activity may not affect plants' operations directly and may have little impact on the number of jobs actually created. It is useful to distinguish this activity from organic growth to understand whether there are differences between the two. Establishment-level data can be used to separate plant-level changes from ownership changes, and separate merger-and-acquisition activity from "organic" growth in firm-level data.

Firm- and establishment-level data may also reach different conclusions about the age of a business entity. Firms can seem to be born from venerable establishments with a stroke of pen when plant mergers occur. As well, if a firm has changed its name or address, some databases may classify it as "new," although its products, workforce and capital remain the same. Based on its plants, it would make sense to classify this firm as a mature incumbent instead of a birth, but without knowing the age of the establishments, researchers must rely on the date of firm's first appearance in the dataset.¹⁰

2.2.2 Regression-to-the-mean

HJM also argue that inferences about firm size and jobs are clouded by regression-to-the-mean: the tendency for firms experiencing transitory shocks to return to their long-run averages, thereby creating a correlation between size and growth. For example, small firms may become

10. On the other hand, the firms may have kept their workers and capital, but changed management, joined another industry, or otherwise changed substantially. In such cases, classifying a firm according to its last change may be more accurate than categorizing it according to the age of its oldest plant. For instance, Nokia, Finland's mobile telephone maker, started as a lumber mill in 1896, and did not move exclusively into telecommunications until the 1990s. In Nokia's case, its oldest establishment is a less useful guide to its current characteristics than current firm-level changes.

temporarily large in response to a greater-than-normal one-off shock and then fall back to their typical size; large firms that lay off workers or cut shifts will appear small in one year, but grow the next. Depending on the method of grouping firms in size classes, such shocks can make large firms look like destroyers, and small firms, creators, of jobs. In many cases, a more accurate interpretation is that small firms create and destroy a large proportion of all jobs, but over the long-term, are responsible for few net new jobs.

Birch's base-year methodology for determining firm size (classifying firms based on their size in period $t-1$ and measuring their growth between $t-1$ and t relative to their size in $t-1$) is particularly susceptible to these effects. An alternative is to use the end year (t) to classify firms, but this produces the opposite bias: the small firms in the example given above will appear particularly liable to destroy jobs, and the large firms, to create them.

Davis et al. (1993) and HJM use an approach that has become standard in the literature—the “current year”—referred to here as the “average-year” method. Firms are classified according to their average size over $t-1$ and t , which is also used as the growth rate denominator.¹¹ This method reduces, though it does not eliminate, regression-to-the-mean effects.

A problem it creates, however, is a tendency to misclassify permanent shocks. A small firm that permanently expands between two periods would be classified correctly by the base-year methodology as a rapidly growing small firm, but incorrectly as a rapidly growing large firm when the average-year methodology is used. Depending on how the size classes are constructed, the average-year method may get it right if the resulting average size is still small, but wrong, if it is large.

The literature is unclear about whether researchers should control for regression-to-the mean. In the strictest interpretation of Gibrat's Law, shocks are uncorrelated over time. Studies testing the properties of shocks in the context of Gibrat usually find autocorrelation, with most finding negative autocorrelation (Petrunia, 2008 and some of the industries in Audretsch et al., 2004). Negative autocorrelation is consistent with regression-to-the-mean effects, which these studies consider to violate the independence of employment growth and firm size.¹²

However, influences on firm growth are commonly separated into temporary and permanent shocks. It is sometimes implicitly argued that correction for the former is required, but that the latter are what is of interest in growth studies, and therefore, should not be excluded from the data. Of course, the validity of this argument depends on short-duration movements really being transitory, an issue that few studies have investigated.

2.2.3 Firm age effect

Neumark et al. (2008) use both establishment and firm data, and treat carefully the regression-to-the-mean issue. They still find a negative relationship between growth rates and firm size in the U.S. private sector, though the importance of small firms in their study turns out to be an order of magnitude less than Birch's claims.

HJM emphasize a third issue that is generally neglected in earlier studies: the role of firm age. Firms are usually born small, and new firms either grow rapidly to become larger, or else they disappear. The data record a large number of net job growth by small start-ups or young firms; but observations on the growth of slower growing small firms are often cut short by their death and thus the sample selected is either not fully representative (if just continuers are considered)

11. When classifying firms according to their base-year size, HJM find a negative relationship between firm size and net growth at the lower end of the firm size distribution (fewer than 500 employees). This relationship disappears almost completely when firms are sorted by their average size.

12. Other studies find positive correlations that are consistent with permanent shocks, or at least temporary shocks lasting longer than a year.

or the observed growth rates are censored downward since observations of decline below zero cannot be measured. The result is the appearance of a negative relationship between size and growth, when growth is actually more a function of characteristics that are summarized more robustly by age. Unfortunately, most data products contain little information on firm age, making a richer characterization of firm growth difficult.¹³

2.3 Firm size distributions

Another body of literature examines the size distribution of firms. Different growth processes generate different distributions of firm size. Gabaix (1999; 2009) shows that growth rates that are independent of size generate a power law distribution, specifically, a Zipf distribution.

Indeed, most studies find that the distribution of firms by size is skewed to the right and that its upper tail follows such a "power law." The tail distribution of firms is often said to follow "Zipf's Law" or to have a "Pareto" distribution. Zipf's Law states that the probability of a firm of size s being larger than a given size S is $\Pr(s \geq S) = aS^{-\xi}$, where a is a scaling parameter and ξ is the Zipf exponent. If Zipf's Law holds, then ξ equals -1.

In general, the Zipf distribution for firm sizes seems to be a good fit for firms larger than a minimum size, with the possible exception of the very largest. Axtell (2001) finds that the universe of U.S. firms (including zero employers, that is, self-employed) had an exponent of -1.06 in 1997, and thus, is well approximated by a Zipf distribution. He states that "the Zipf distribution is an unambiguous target that any empirically accurate theory of the [U.S.] firm must hit," and suggests that the Zipf distribution may describe firm sizes in other countries as well. Fujiwara et. al. (2003) report similar results for European firm employment, although their data consist only of larger firms.

13. Firm age has long been recognized as an important factor even if it appears in only a few empirical studies: see Evans (1987) for an early example. Birch (1981) highlighted the importance of age: "Another distinguishing characteristic of job replacers is their youth Not all small businesses are job creators. The job creators are the relatively few younger ones that start up and expand rapidly in their youth, outgrowing the 'small' designation in the process" (p. 8).



3 The LEAP database

This analysis uses Statistics Canada's Longitudinal Employment Analysis Program (LEAP) database.¹⁴ The LEAP, an administrative database, includes all firms that have some payroll, and therefore, issue at least one "statement of remuneration paid"—a T4-slip. It covers incorporated and unincorporated businesses, but excludes self-employed individuals or partnerships where the participants do not draw salaries. The LEAP data cover the 1983-to-2008 period. Using the information gathered by Statistics Canada's Business Register, the administrative data in LEAP are structured at the level of the "statistical enterprise," which is the level associated with a complete set of financial statements.¹⁵ In the current analysis, this statistical unit is referred to as "the firm."

Although the LEAP database contains information only at the level of the firm, it has several advantages. Because the LEAP pertains to the entire economy, using it avoids the manufacturing bias of many other studies. And whereas many studies rely on data with limited coverage at the smaller end of the firm size scale or impute data for small firms, the LEAP has data on almost all firms.¹⁶

The chief advantage of the LEAP database is that a major form of corporate restructuring, which creates problems in employment growth studies, can be removed. In a longitudinal panel, mergers, acquisitions, spinoffs and divestitures can result in the spurious "creation" and "destruction" of firm identifiers. For example, if a large firm acquires a small firm, the small firm seems to have died, and the consolidated firm might appear in the data with its own identifier, thereby looking like a new firm. In reality, the acquisition resulted in neither a firm birth nor a death. The employees of the large firm still work for the same large firm, and most employees of the small firm, for the same establishments.

These false births and deaths are removed from LEAP through "labour-tracking." Clusters of employees of appearing and disappearing firms are compared with the clusters of employees of other firms in the previous year (for appearing firms) or the following year (for disappearing firms). If substantial portions of the employees of appearing and disappearing firms are found previously or subsequently in another firm, a connection is made between those two firms, and the structure of the firm in year t is applied to the data of both firms in year $t-1$ and in all preceding years. Thus, employment growth between any two years will be attributed to firms of different size classes based on the structure of the firm in the final year (year t).¹⁷

14. See Baldwin, et al. (1992) for a description of the construction of the database.

15. "The enterprise, as a statistical unit, is defined as the organisational unit of a business that directs and controls the allocation of resources relating to its domestic operations, and for which consolidated financial and balance sheet accounts are maintained from which international transactions, an international investment position and a consolidated financial position for the unit can be derived." <http://www.statcan.gc.ca/concepts/definitions/ent-eng.htm>.

16. Its coverage of the self-employed is not comprehensive, including only the self-employed with workers.

17. This process chooses the structure in the second year as the point of comparison. While this choice is arbitrary, so, too, would be the choice of the structure in the preceding year.

The LEAP database is produced in vintages. When an additional year of data becomes available, a new longitudinal file is created, but the previous longitudinal files are also retained. Microdata files are available for each vintage starting in 1999 up to the most recent, 2008. LEAP 2008 tracks firm employment from 1983 to 2008; 2007 tracks employment from 1983 to 2007, and so on back to 1999, which track employment from 1983 to 1999. Each vintage holds the structure of the firm constant, based on information for the most recent year available. For example, 2008 aggregates all earlier annual data based on the firm's structure in 2008; that is, firms that merged in 2008 are consolidated in all previous years in the 2008 vintage.

This practice makes it possible to measure only "organic" employment growth, and ignore employment changes caused by merger and acquisition activity. For example, between period $t-1$ and t , firm B acquires firm A's 100 workers (growth due to mergers and acquisitions) and also adds 25 employees to the combined entity (organic growth) (Appendix 1). Had the LEAP file not been altered as described above, it would record a sharp jump in firm B's employment in year t (the 100 employees from firm A, plus 25 new positions), and a 100-employee decrease in firm A's employment due to its (false) death.¹⁸

Instead, the change in employment recorded between $t-1$ and t removes the acquisition transaction and reflects only the employment growth that occurred "within" the now-larger firm B, which rises from 700 employees (B's original 600 plus 100 from firm A) in year $t-1$ to 725 employees in year t . This larger firm existed in both year $t-1$ and year t . Classifying growth over the previous year to the consolidation recognizes only that this growth is attributed to the newly created entity, since it was responsible for operations over this period. Thus, year-on-year, LEAP abstracts almost completely from plants changing hands, recording only "organic" growth in the overall entity defined in terms of its structure in year t .

The fact that each vintage pushes the market structure of its latest year back in time creates an extra challenge. The practice risks misclassifying the sources of growth by firm size for every year except the last two. The structure of a firm in $t-2$ is clearly different than the longitudinal structure, and it is not the case that the firm in year t was in any sense responsible for the consolidated operations in year $t-2$.

In the t vintage of the data in the previous example, the records for firms A and B are merged into one unique record for all years preceding year t , adding up the employment of both firms. This synthetic record is then linked with firm B's record in year t . Firm A disappears completely from the t vintage, and its employees are treated as though they were always part of firm B—from their first appearance to the end of the vintage.

To deal with the potential for misclassification in earlier years, this analysis uses only the last two years of each LEAP vintage from 1999 to 2008. The result is a series of ten cross-sections, pooled to estimate the distribution of firm sizes, firm ages, and employment growth rates. This reduces the comparisons to one-year periods, with the firm structure held constant for the second year of each cross-section. This approach should yield accurate inference, provided that the distribution is stationary and the sample period is not characterized by large abnormal shocks. As will be explained later, the longitudinal data prior to 1998 back to 1983 are used only to calculate firm age.

3.1 Measures of employment

No measure of employment fits all purposes, and every measure has limitations. Estimates of the number of workers tend to be either a count employment at a fixed date each year or a

18. False in the sense that firm A's economic activity continues, even if its identifiers (business number, name, etc.) do not.

calculation of the average level of employment during the year using infra-annual estimates.¹⁹ The Longitudinal Business Database used by HJM measures employment directly once a year as the number of workers at each establishment in the week that includes March 12. While this measure is consistent, late March start-ups are identified as firm births the following year although they are almost a year old. Also, employees with more than one job can be double-counted. And finally, seasonal patterns that vary across firms and industries are missed.

The LEAP, by contrast, contains annual payroll information (total payroll and total number of workers who received a T4 slip over the calendar year). Based on these data, employment levels are inferred through two measures: Average Labour Unit (ALU) and Individual Labour Unit (ILU).

The ALU divides each firm's annual payroll by the average annual earnings of a representative worker in the same industry, province and firm size class²⁰ because earnings have been found to vary along these dimensions (Statistics Canada, 2011; Baldwin, 1998; Drolet and Morissette, 1998). The average annual earnings for the various industry, province and firm size class combinations are determined from Statistics Canada's monthly Survey of Employment, Payrolls and Hours. For a detailed description of the ALU calculation, see Lafrance and Leung (2010).

The ALU can vary along two dimensions—number of workers, and average remuneration, which differs across firms because of longevity of employment over a year, hours worked and wages levels. Firms with a small ALU may, indeed, have fewer employees, and firms with a larger ALU may have more. However, firms may appear small because they pay wages below the industry average, or large because they pay more. The ALU, therefore, takes several dimensions of job 'quality' into account.²¹

As a result, compared with simple job count measures, the ALU may enlarge differences between small and large firms. Whether this is undesirable depends on the objective—to study raw job creation or quality-corrected job creation. Because large firms are not numerous and have many employees, they dictate the average annual earnings in their provincial industry and size class. Average annual earnings for small firms are defined over a much larger population of firms, which implies more variation.

The ALU calculation considers only four firm size classes: 200 employees or more, 50 to 199, 20 to 49, and fewer than 20. Although the smallest category groups all firms with fewer than 20 employees, in many industries, firms with fewer than 5 employees offer lower earnings than do those with 5 to 19 employees (Statistics Canada, 2011). Many small firms, therefore, have an ALU value below 1. This could also occur if firms are active for only part of the year, which is more likely for start-ups created at the end of a calendar year.²²

To check the robustness of the results, the Individual Labour Unit (ILU) is also used. The ILU is closer to a head count—every individual who receives at least one T4 slip in a given year is assigned 1.0 ILU. If individuals worked for different firms during the year, their 1.0 ILU is split proportionately across firms according to the share of their total annual payroll earned in each. For example, if someone earned \$10,000 in firm A, and \$40,000 in firm B, then firm A is assigned 0.2 ILU, and firm B, 0.8 ILU.

19. Measures of employment in the Productivity Accounts also control for the intensity of work by counting hours worked.

20. If a firm operates in multiple provinces, its national ALUs correspond to the sum of its provincial ALUs. Also, the same size class is used in all provinces because firm size is considered to be a national characteristic. However, the dominant industry is allowed to change across provinces for a given firm.

21. Baldwin (1998) examines job growth, taking job quality into account based on average wage.

22. Because the precise date when new firms started is not known, part-year firms cannot be identified with certainty. For births, it is difficult to adjust for size because many firms do not survive in the next year. As a result, births might appear smaller than they actually were.

The ILU ignores differences in the quality of jobs. Because it counts the number of T4 slips issued, regardless of the number of hours actually worked, it can treat a four-hour part-time worker in the same way as a forty-hour full-time staffer. Firms that hire many part-time or casual workers will appear to have a large number of jobs, even if the total number of hours worked is small.

On average, a firm's ILU value in the LEAP file is 3.0 times greater than its ALU employment (Appendix, Table 9). This average is heavily influenced by the employment patterns of tiny firms (less than one ALU), for which the average ILU value is 5.7 times greater than its ALU value. By contrast, the ILU for firms with 20 or more employees is, on average, only 1.1 times greater than their ALU. The wider disparity in the measures at the lower end of the size scale is likely a function of the lower wages earned by the self-employed, start-up workers, and employees of part-year firms.

Considering how the two measures are obtained, it is not surprising that more ILUs than ALUs are enumerated annually in the business sector (Appendix, Chart 13). For example, in 2008, ILUs numbered 13.5 million, compared with 11.5 million ALUs. The magnitude of gross job creation and gross job destruction is also larger based on ILUs than ALUs (Appendix, Chart 14). However, the levels of net job growth are similar. The dynamics of net job growth tend to follow the same trend using the ALU and ILU—when one rises, so does the other. Some evidence suggests that the measures respond differently over the business cycle, with increases in ILUs leading ALUs. This could be the result of firms responding to increased demand by hiring (lower-wage) part-time employees, and then making them permanent if demand persists.

3.2 Defining age

Because the LEAP employment data go back to 1983, it is possible to obtain a measure of firm age. HJM define firm age in the following way: "... when a new firm ID arises for whatever reason, firm age is based upon the age of the oldest establishment that the firm owns in the first year the new firm ID is observed. The firm is then allowed to age naturally ... regardless of mergers or acquisitions and as long as firm ownership and control do not change" (p. 13).

The definition of firm age in this analysis is similar. For every firm ID in the study, the first year it is observed in the 1999-to-2008 period is identified, and the corresponding LEAP vintage is used (for example, the 2000 vintage if the firm first appeared in 2000). In this vintage, the earliest year in which the firm had employees becomes its birth year.^{23,24} Thereafter, the firm is aged naturally.²⁵

An age of zero corresponds to a birth. A birth is identified when all the constituents of the new firm have no employment history. This definition fits well with that of HJM and the definition of enterprise birth proposed by the OECD and Eurostat: "A birth amounts to the creation of a

23. When a new firm ID appears, the age assigned to it is not necessarily zero. For example, two firms might merge to create a legitimately new firm with a different ownership and a new line of products. Because employment in the past can be traced for the constituents of the new firm, its age is not zero. This situation is treated in a similar way by HJM: the age of the new firm will not be zero in its first year, if it acquired establishments that already existed (whose age is not zero).

24. For firms present in the first year of the sample used here (1999), it is not possible to obtain the age the first time the firm ID ever appeared, because microdata are not available before 1999. The distribution of firm age is only slightly different for firms that first appeared in the sample in 1999, compared with those that first appeared later—fewer firms are assigned an age of zero or 1 when they first appear, and more are assigned an age of 2 or more. The greatest difference is for age zero (92% of firms present in 1999, compared with 96% for firms that appeared later). However, this does not seem to introduce a bias: when cross-sections for firms that first appeared in 1999 are excluded from the regressions, the conclusions are the same.

25. If firms have no employees for a year or more, and then resume having employees, they are still aged during their absence from the employer population. The vast majority of these firms register sales even if they have no employees.

combination of production factors with the restriction that no other enterprises are involved in the event. Births do not include entries into the population due to mergers, break-ups, split-off or restructuring of a set of enterprises" (Eurostat/OECD, 2007).

In the 1999 vintage of LEAP, the birth year of firms can be identified back to 1984 (15 years old in 1999), but the exact age of older firms is not known. For consistency, a separate age category, "16 or older," is created for all subsequent vintages. The age variable used by HJM is also right-censored at 15, with a group for firms aged 16 or more years.

3.3 Empirical framework: Growth

For the main measures of firm size and employment growth, the "current" or "average" year method of HJM is adopted here. The average size of firm i is defined as $X_{it} = \frac{E_{it} + E_{it-1}}{2}$, and

employment growth is defined as $g_{it} = \frac{E_{it} - E_{it-1}}{X_{it}}$, where g_{it} is firm i 's growth rate from period $t-1$

to period t , and E is employment. Note that the results on rates are not presented as percentage changes in the text and tables: a value of 1 corresponds to 100%, a value of 0.05 to 5%.

This formulation, which is a standard definition of employment growth in the empirical literature on Gibrat's Law, has a number of desirable properties. First, it has the advantage of symmetry, in that it is invariant to the choice of a base year.²⁶ Second, it yields similar results to log growth rates, at least for firms increasing or reducing employment by up to 50% (by log/average measures). The chief difference from log growth is that the "average-year" method is linear and bounded by -2 and 2 at the extreme ends. The "average-year" method assigns births and deaths a growth rate of 2 and -2, respectively, by construction. As a result, births and deaths constitute a distinct subset that is not directly comparable to other observations and is best analyzed separately. In addition, average employment in year t is the average of $(0, E_{it})$ for births and $(E_{it-1}, 0)$ for dying firms, meaning that the denominator for these firms arbitrarily corresponds to half their actual size upon entry and exit.

The third advantage of the "average-year" method is that this measure is easily aggregated. For example, defining the growth rate in aggregate employment as

$$g_t = \sum_s X_s g_{st} = \sum_s X_s \sum_{i \in s} \frac{X_{it}}{X_s} g_{it} \quad (1)$$

where $s = 1, \dots, S$ is some grouping of observations according to some identifiable characteristics, such as size or age. X_s and g_s are defined using the group's total employment in $t-1, t$. The X_s 's are a measure of the group's relative size importance in the economy, and g_s is the group's growth rate.

3.4 Empirical framework: Size classification

The average-year methodology of HJM differs from Birch's base-year methodology in two respects: the denominator used in the growth rate calculation, and the method used to assign a firm to a size class. Okolie (2004) demonstrated that, when calculating net and gross employment growth, the method of classifying firms into size categories matters more than the denominator used. For this reason, only one denominator is used in this study, namely the average size in $t-1$ and t . To determine whether regression-to-the-mean is present in the LEAP

26. See Tornqvist, et al. (1985) for details on the desirable properties of growth measures.

data, the results obtained when firms are classified according to their base-year size (size in $t-1$) are compared with results using their average size (average of size in $t-1$ and t).

The average-size classification mitigates, although it does not eliminate, regression-to-the-mean problems. It takes care of short-term regression-to-the-mean of small or medium magnitude, but major transitory shocks or shocks lasting more than one period may not be completely averaged out. However, as previously discussed, it is not clear that persistent effects of shocks should be defined as transitory and eliminated. The method used here also tends to misclassify enduring shocks. For instance, a small firm that permanently adds jobs may become, on average, a medium-size firm. But its correct classification is that of a small job creator. In this case, the base-year methodology produces the right answer, and the average size does not.

3.5 Econometric framework

The present analysis is restricted to the business sector. This is done by excluding the cross-sections classified in NAICS 91 (public administration), 61 (education services) and 62 (health care and social assistance).²⁷

To get a picture of the statistical relationship of employment growth rates to firm age and size, the following regression model is specified:

$$g_{it} = \Gamma'Z + \sum_{i=1}^N \left(\sum_{s=1}^S \beta_s \cdot 1 \cdot [i \in s] \right) + \sum_{a=0}^A \gamma_a \cdot 1 \cdot [i \in a] + \xi_{it} \quad (2)$$

where $s = 1, \dots, S$ are firm size classes, $a = 1, \dots, A$ are firm age classes, and $1[\cdot]$ is an indicator function that equals 1 if firm $i = 1, \dots, N$ belongs to the given size/age class. The vector Z contains control variables (dummies for year and the 2-digit NAICS industry). Omitted categories are year 2003 and manufacturing (NAICS 33).

Five different variations of the regression are employed, using both the ALU and ILU measure of employment:


- Size controls only, using base size classification
- Size controls only, using average size classification
- Age controls only
- Size and age controls, using base size classification
- Size and age controls, using average size classification

In each specification, industry and year dummies are included. The coefficients on the dummy variables (indicator functions) are, in effect, conditional means of the growth rates of particular groups. Of interest is whether the β 's or γ 's differ across firm size and age groups, respectively.

In the regressions, the cross-sectional observations with a growth rate equal to 2 (entry), but for which age is not zero, are excluded.²⁸ Even though they are not numerous, these observations could bias conclusions about the effect of firm age.

27. Canada's education and health sectors are heavily financed by the public sector. Unlike firms in the private sector, the firms in the excluded sectors are not necessarily profit-maximizing.

28. HJM's job creation and destruction calculation excludes firms that re-enter after a temporary absence. Because two-year cross-sections are being used, these absences cannot be directly identified. By imposing the condition that age is non-zero, these re-entries are excluded; but firms that first appear with an age greater than zero are also excluded (4% of firm appearances).



4 Results

Firms with less than one ALU or ILU are designated “tiny” firms. “Small” firms are defined as having fewer than 500 employees, and “large” firms, more than 500 employees (using either the ALU or the ILU measure). “Start-ups” are aged zero; “young” firms are less than 10 years old; and “mature” firms are 10 or more years old.

4.1 Descriptive statistics

4.1.1 Distribution of firms and employment across size and age classes

As of 2008, the business sector comprised more than 1 million incorporated and unincorporated firms with employees (Table 1). Almost 4 in 10 (38%) of these firms had less than one ALU, which may indicate very few workers (one self-employed or working part-year only) and/or especially low wages compared with other small firms. Indeed, 25% of these tiny firms had exactly one ILU employee, and almost two-thirds had one or less. However, firms that start up late in the year with employees who earned the bulk of their incomes with other employers could also yield this result. Close to half (48%) the firms—slightly more than half a million—had more than one but fewer than 10 ALUs. Another 11% (around 115,000) had 10 to 49 employees. Firms with 50 or more employees represented 2.7% of the total, and large firms with 500 or more employees, only 0.2%.

Table 1
Business sector firms, by age and size, 2008

Firm size	Age					Total	Age					Total
	Less than 1 year	1 to 3 years	4 to 9 years	10 to 15 years	16 years or older		Less than 1 year	1 to 3 years	4 to 9 years	10 to 15 years	16 years or older	
	number						percent					
Average labour units												
0 to less than 1	86,066	117,403	93,860	49,038	50,799	397,166	8.2	11.2	9.0	4.7	4.9	38.0
1 to less than 10	36,257	117,203	140,259	86,152	125,384	505,255	3.5	11.2	13.4	8.2	12.0	48.3
10 to less than 50	1,784	13,698	26,584	21,598	51,618	115,282	0.2	1.3	2.5	2.1	4.9	11.0
50 to less than 100	107	1,376	3,139	2,780	9,044	16,446	0.0	0.1	0.3	0.3	0.9	1.6
100 to less than 500	x	x	1,604	1,570	6,389	10,327	x	x	0.2	0.2	0.6	1.0
500 or more	x	x	203	207	1,350	1,842	x	x	0.0	0.0	0.1	0.2
Total	124,274	250,466	265,649	161,345	244,584	1,046,318	11.9	23.9	25.4	15.4	23.4	100.0
Individual labour units												
0 to less than 1	35,321	31,931	25,863	12,916	11,931	117,962	3.4	3.1	2.5	1.2	1.1	11.3
1 to less than 10	84,783	195,630	199,554	117,332	156,414	753,713	8.1	18.7	19.1	11.2	14.9	72.0
10 to less than 50	3,858	20,120	34,311	25,875	58,225	142,389	0.4	1.9	3.3	2.5	5.6	13.6
50 to less than 100	210	1,794	3,747	3,175	9,511	18,437	0.0	0.2	0.4	0.3	0.9	1.8
100 to less than 500	x	x	1,937	1,800	7,022	11,750	x	x	0.2	0.2	0.7	1.1
500 or more	x	x	237	247	1,481	2,067	x	x	0.0	0.0	0.1	0.2
Total	124,274	250,466	265,649	161,345	244,584	1,046,318	11.9	23.9	25.4	15.4	23.4	100.0

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

About 12% of firms were start-ups, and another 49% were less than 10 years old. The remaining 39% were mature firms. Small firms were disproportionately young (65%), and large firms, predominately old (85%) (Table 1).

Private sector jobs in Canada totalled 11.5 million in 2008, up from 10.2 million in 1999. Firms with fewer than 100 employees accounted for roughly half of private sector employment.

Based on the ALU measure, tiny firms make up less than 2% of all employment, and large firms, 35% (Table 2). Based on the ILU measure, small firms' share of employment is slightly larger, and large firms', correspondingly lower. This, again, suggests that small firms generally have more low-wage and part-time jobs.

Table 2
Employment, by firm age and firm size, 2008

Firm size	Age					Total	Age					Total
	Less than 1 year	1 to 3 years	4 to 9 years	10 to 15 years	16 years or older		Less than 1 year	1 to 3 years	4 to 9 years	10 to 15 years	16 years or older	
	number						percent					
Average labour units												
0 to less than 1	32,009	54,663	44,953	23,887	25,373	180,886	0.3	0.5	0.4	0.2	0.2	1.6
1 to less than 10	92,069	349,804	467,731	301,931	487,480	1,699,016	0.8	3.0	4.1	2.6	4.2	14.7
10 to less than 50	33,301	271,029	537,679	444,221	1,126,680	2,412,909	0.3	2.3	4.7	3.8	9.8	20.9
50 to less than 100	7,259	92,542	213,224	189,460	620,804	1,123,289	0.1	0.8	1.8	1.6	5.4	9.7
100 to less than 500	x	x	294,961	287,109	1,224,523	1,949,343	x	x	2.6	2.5	10.6	16.9
500 or more	x	x	349,516	478,259	3,270,052	4,179,515	x	x	3.0	4.1	28.3	36.2
Total	180,424	976,691	1,908,064	1,724,868	6,754,912	11,544,959	1.6	8.5	16.5	14.9	58.5	100.0
Individual labour units												
0 to less than 1	17,236	16,478	13,249	6,594	6,230	59,788	0.1	0.1	0.1	0.0	0.0	0.4
1 to less than 10	180,563	514,217	601,927	373,941	569,825	2,240,472	1.3	3.8	4.4	2.8	4.2	16.5
10 to less than 50	70,318	391,261	679,359	527,565	1,256,411	2,924,915	0.5	2.9	5.0	3.9	9.3	21.6
50 to less than 100	14,192	121,818	255,519	218,392	659,514	1,269,435	0.1	0.9	1.9	1.6	4.9	9.4
100 to less than 500	x	x	353,099	337,775	1,342,713	2,216,446	0.1	x	2.6	2.5	9.9	16.4
500 or more	x	x	421,964	612,301	3,680,688	4,831,523	0.1	x	3.1	4.5	27.2	35.7
Total	307,729	1,317,783	2,325,117	2,076,569	7,515,381	13,542,578	2.3	9.7	17.2	15.3	55.5	100.0


Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

4.1.2 Mean growth rate by size and age

Employment growth rates by firm size classes depend on whether the base-size or the average-size method is used to classify firms (Charts 1, 2 and 3). With the base-size classification, the average growth rate is negative for all size classes except tiny firms.²⁹ The opposite occurs with the average-size classification: except for tiny firms, all size classes have a positive growth rate. These divergent results illustrate the importance of the sizing method used. The different sign for the rate of growth indicates how much regression to the mean occurs—and that firms that are likely to change direction are more likely to have larger negative than positive growth, perhaps because decline possibilities are less likely to be bounded.

Chart 1

Average annual employment growth rate, 1999 to 2008 — By firm size, with base size classification




Note: The average size is used in the denominator of the growth rate in all cases.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

29. When classifying firms according to their base-year size, but defining the growth rate as $(X_t - X_{t-1})/X_{t-1}$ instead of $(X_t - X_{t-1})/\text{avg}(X_t, X_{t-1})$, the average growth rates are positive for all size classes. They are also declining in size. Only continuing firms were considered in this exercise, because the standard growth rate is not well defined for entries and exits.

Chart 2

Average annual employment growth rate, 1999 to 2008 — By firm size, with average size classification



Note: The average size is used in the denominator of the growth rate in all cases.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Chart 3

Average annual employment growth rate, 1999 to 2008 — By firm age



Note: The average size is used in the denominator of the growth rate in all cases.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

With the average-size classification, the rate of growth rises over the smallest firm size classes, is relatively constant for the middle size classes, and declines in the largest size classes. This pattern prevails whether employment is measured with the ALU or the ILU, and whether all firms or just continuing are examined.

With the base-size classification, the rate of employment growth is larger at the low end of the firm size distribution, namely, among firms with less than 1 ALU or ILU. However, with the average-size classification, the rate of growth is lower for these micro-firms. In either case, micro-firms are probably sufficiently different from the rest of the population of firms that conclusions about the superior growth of small firms should probably not rest on this category.

The specific size categories for this analysis were chosen to provide a picture that allows us to evaluate whether there are monotonic trends from small to large groups that reveal a uniform superiority of ever smaller firms. It is difficult to see a particular strong trend in average growth rates that suggests its existence using the average size classification. When the ILU measure is used for all firms, middle sized firms grow faster than both small and large firms. When the ALU measure is used, the pattern is somewhat irregular above 1 ALU, though the largest firms have slower growth rates than all but micro firms.

It should be recognized that differences across size classes are rarely investigated in this amount of detail. Indeed, very aggregative groupings are often used that arbitrarily divide the universe of firms into three groups: less than 100 employees, 100 to less than 500 employees, and 500 employees or more. Using these broad size classes, smaller firms are found to have higher growth rates than middle sized firms and the largest firms (Table 3). Aggregations such as these hide whether it is all small firms that have superior performance or only a subset, which the more detailed categories used here allow us to examine. But all of these categories (the very aggregate ones and the more detailed ones) fail to take into account differences that might be expected in growth rates across groupings that are due to compositional effects, that is, whether small firms are found more in some industries that might be experiencing more rapid growth over the period studies, or whether time periods chosen affect the relative importance of small versus large firms or whether differences in age is an important factor and affects differences found across size classes.

Table 3
Average growth rate using the average size classification, by broad firm size classes, 1999 to 2008

Firm size	Average labour units	Individual labour units
	average growth rate	
0 to less than 100 (reference group)	0.0461	0.0382
100 to less than 500	0.0272 *	0.0203 *
500 or more	0.0209 *	0.0132 *

* indicates statistically different from smallest size class at the 0.05 level.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Chart 3, which presents the average employment growth rate by age, demonstrates the importance of a firm's early years. Many start-ups do not survive to the next year, but those that do survive grow—the average growth rate for continuing one-year-old firms is positive.³⁰ Thereafter, average employment growth rates are, on average, negative. The growth rate is less for firms overall than for continuing firms, because exits reduce the average substantially. Accounting for firm age, therefore, provides more information about employment growth than does firm size alone.

30. One-year-old continuers form the only group where growth rates are not distributed symmetrically around zero. They are skewed to positive growth. This is due to size being measured imperfectly in the birth year when a firm is active only part of the year.

The next section make use of multivariate analysis to consider what happens to differences in growth rates across size classes once differences in industry, time period and age characteristics are considered.

4.2 Regression results – Conditional means

The regression results are presented in Tables 4 to 7. The parameters from each of these regressions are graphed in Charts 4 to 11. To each parameter, the unconditional mean of the growth rates for the excluded size class (at least 1,000 employees) and/or age category (16 years old or more) is added. As noted above, the unconditional mean for the larger size class is negative with the base-size classification, but positive with the average-size classification.

Table 4
Regression results, using average labour units — All firms

Variables	Base size		Average size		Firm age only		Base size with firm age controls		Average size with firm age controls	
	estimate	standard error	estimate	standard error	estimate	standard error	estimate	standard error	estimate	standard error
Firm size¹ (ALU)										
< 1	0.2378 ***	0.0036	-0.0958 ***	0.0033	-0.3604 ***	0.0035	-0.6484 ***	0.0032
1 to <5	-0.1631 ***	0.0035	-0.0213 ***	0.0033	-0.1414 ***	0.0035	-0.1374 ***	0.0031
5 to < 10	-0.0830 ***	0.0035	-0.0075 *	0.0033	-0.0564 ***	0.0035	-0.0403 ***	0.0031
10 to < 20	-0.0516 ***	0.0035	0.0004	0.0033	-0.0267 ***	0.0035	-0.0183 ***	0.0032
20 to < 50	-0.0377 ***	0.0035	0.0040	0.0033	-0.0136 ***	0.0035	-0.0006	0.0032
50 to < 100	-0.0264 ***	0.0036	0.0096 **	0.0034	-0.0058	0.0036	0.0094 **	0.0032
100 to < 250	-0.0221 ***	0.0037	0.0092 **	0.0035	-0.0107 **	0.0037	0.0042	0.0033
250 to < 500	-0.0297 ***	0.0046	0.0072	0.0044	-0.0261 ***	0.0046	-0.0027	0.0041
500 to < 1,000	-0.0151 **	0.0053	0.0014	0.0051	-0.0145 **	0.0053	-0.0084	0.0048
Firm age² (years)										
0	2.1857 ***	0.0005	2.4043 ***	0.0008	2.5415 ***	0.0007
1	0.0448 ***	0.0012	0.1877 ***	0.0012	0.2895 ***	0.0012
2	-0.2710 ***	0.0011	-0.1764 ***	0.0011	-0.0871 ***	0.0010
3	-0.2129 ***	0.0011	-0.1318 ***	0.0011	-0.0566 ***	0.0011
4	-0.1675 ***	0.0012	-0.0960 ***	0.0011	-0.0312 ***	0.0011
5	-0.1423 ***	0.0012	-0.0784 ***	0.0012	-0.0214 ***	0.0011
6	-0.1235 ***	0.0012	-0.0650 ***	0.0012	-0.0133 ***	0.0012
7	-0.1032 ***	0.0012	-0.0496 ***	0.0012	-0.0029 *	0.0012
8	-0.0869 ***	0.0013	-0.0375 ***	0.0013	0.0050 ***	0.0012
9	-0.0777 ***	0.0013	-0.0320 ***	0.0013	0.0066 ***	0.0013
10	-0.0691 ***	0.0014	-0.0269 ***	0.0013	0.0088 ***	0.0013
11	-0.0599 ***	0.0014	-0.0218 ***	0.0014	0.0100 ***	0.0013
12	-0.0459 ***	0.0014	-0.0117 ***	0.0014	0.0162 ***	0.0014
13	-0.0396 ***	0.0015	-0.0097 ***	0.0015	0.0147 ***	0.0014
14	-0.0323 ***	0.0015	-0.0050 ***	0.0015	0.0172 ***	0.0015
15	-0.0256 ***	0.0016	-0.0004	0.0015	0.0197 ***	0.0015

	Base size	Average size	Firm age only	Base size with firm age controls	Average size with firm age controls
Diagnostic statistics					
Number of observations (millions)	10.5	10.5	10.5	10.5	10.5
R-squared	0.0461	0.0190	0.4612	0.4726	0.5155

* significantly different from reference group at the 0.05 level.

** significantly different from reference group at the 0.001 level.

*** significantly different from reference group at the 0.0001 level.

1. Reference group is 1,000 or more, except for "Firm age only."

2. Reference group is 16 or older, except for "Base size" and "Average size."

Note: ALU = Average Labour Units. Standard errors are correct for heteroscedasticity. Year and two-digit North American Industry Classification System (NAICS) industry controls are also included in these regressions.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Table 5
Regression results, using average labour units — Continuing firms

	Base size		Average size		Firm age only		Base size with firm age controls		Average size with firm age controls	
	estimate	standard error	estimate	standard error	estimate	standard error	estimate	standard error	estimate	standard error
Variables										
Firm size¹ (ALU)										
< 1	0.1733 ***	0.0032	-0.0387 ***	0.0030	0.0855 ***	0.0032	-0.1482 ***	0.0030
1 to <5	-0.0113 ***	0.0031	-0.0044	0.0030	-0.0406 ***	0.0032	-0.0643 ***	0.0030
5 to < 10	-0.0142 ***	0.0031	0.0053	0.0030	-0.0246 ***	0.0032	-0.0271 ***	0.0030
10 to < 20	-0.0047	0.0032	0.0092 **	0.0030	-0.0102 **	0.0032	-0.0130 ***	0.0030
20 to < 50	-0.0056	0.0032	0.0098 **	0.0030	-0.0085 **	0.0032	-0.0056	0.0030
50 to < 100	-0.0046	0.0032	0.0109 ***	0.0031	-0.0057	0.0032	0.0012	0.0030
100 to < 250	-0.0063	0.0033	0.0104 ***	0.0031	-0.0082 *	0.0033	0.0024	0.0031
250 to < 500	-0.0124 **	0.0040	0.0096 *	0.0039	-0.0150 ***	0.0040	0.0021	0.0038
500 to < 1,000	-0.0030	0.0046	0.0013	0.0045	-0.0064	0.0047	-0.0043	0.0045
Firm age² (years)										
1	0.4518 ***	0.0009	0.4045 ***	0.0009	0.4983 ***	0.0009
2	0.0197 ***	0.0007	-0.0020 **	0.0007	0.0534 ***	0.0007
3	0.0109 ***	0.0007	-0.0068 ***	0.0008	0.0405 ***	0.0008
4	0.0125 ***	0.0008	-0.0024 **	0.0008	0.0389 ***	0.0008
5	0.0072 ***	0.0008	-0.0058 ***	0.0008	0.0311 ***	0.0008
6	0.0062 ***	0.0008	-0.0054 ***	0.0008	0.0282 ***	0.0008
7	0.0063 ***	0.0008	-0.0041 ***	0.0008	0.0267 ***	0.0008
8	0.0057 ***	0.0008	-0.0040 ***	0.0008	0.0247 ***	0.0008
9	0.0046 ***	0.0009	-0.0042 ***	0.0009	0.0222 ***	0.0009
10	0.0031 ***	0.0009	-0.0049 ***	0.0009	0.0196 ***	0.0009
11	0.0035 ***	0.0009	-0.0035 ***	0.0009	0.0184 ***	0.0009
12	0.0052 ***	0.0009	-0.0011	0.0009	0.0186 ***	0.0009
13	0.0014	0.0010	-0.0040 ***	0.0010	0.0134 ***	0.0010
14	0.0013	0.0010	-0.0038 ***	0.0010	0.0123 ***	0.0010
15	0.0034 **	0.0010	-0.0013	0.0010	0.0135 ***	0.0010

	Base size	Average size	Firm age only	Base size with firm age controls	Average size with firm age controls
Diagnostic statistics					
Number of observations (millions)	8.1	8.1	8.1	8.1	8.1
R-squared	0.0298	0.0061	0.0746	0.0844	0.0826

* significantly different from reference group at the 0.05 level.

** significantly different from reference group at the 0.001 level.

*** significantly different from reference group at the 0.0001 level.

1. Reference group is 1,000 or more, except for "Firm age only."

2. Reference group is 16 or older, except for "Base size" and "Average size."

Note: ALU = Average Labour Units. Standard errors are correct for heteroscedasticity. Year and two-digit North American Industry Classification System (NAICS) industry controls are also included in these regressions.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Table 6
Regression results, using individual labour units — All firms

Variables	Base size		Average size		Firm age only		Base size with firm age controls		Average size with firm age controls	
	estimate	standard error	estimate	standard error	estimate	standard error	estimate	standard error	estimate	standard error
Firm size¹ (ILU)										
< 1	0.8454 ***	0.0034	-0.1250 ***	0.0032	-0.4115 ***	0.0035	-1.0210 ***	0.0029
1 to <5	-0.2304 ***	0.0032	-0.0358 ***	0.0030	-0.1863 ***	0.0032	-0.2026 ***	0.0028
5 to < 10	-0.1178 ***	0.0032	-0.0130 ***	0.0030	-0.0752 ***	0.0032	-0.0613 ***	0.0028
10 to < 20	-0.0868 ***	0.0032	-0.0021	0.0030	-0.0470 ***	0.0032	-0.0233 ***	0.0028
20 to < 50	-0.0548 ***	0.0032	0.0065 *	0.0030	-0.0204 ***	0.0033	0.0015	0.0028
50 to < 100	-0.0289 ***	0.0033	0.0106 ***	0.0031	-0.0008	0.0033	0.0146 ***	0.0029
100 to < 250	-0.0190 ***	0.0034	0.0120 ***	0.0032	-0.0004	0.0034	0.0111 ***	0.0030
250 to < 500	-0.0149 ***	0.0040	0.0097 **	0.0038	-0.0028	0.0041	0.0067	0.0035
500 to < 1,000	-0.0105 *	0.0048	0.0027	0.0046	-0.0056	0.0049	-0.0029	0.0043
Firm age² (years)										
0	2.1760 ***	0.0004	2.4573 ***	0.0013	2.7311 ***	0.0008
1	-0.0699 ***	0.0011	0.0403 ***	0.0011	0.2079 ***	0.0010
2	-0.2807 ***	0.0010	-0.2153 ***	0.0010	-0.0924 ***	0.0009
3	-0.2213 ***	0.0011	-0.1646 ***	0.0011	-0.0644 ***	0.0009
4	-0.1764 ***	0.0011	-0.1255 ***	0.0011	-0.0425 ***	0.0010
5	-0.1494 ***	0.0011	-0.1035 ***	0.0011	-0.0325 ***	0.0010
6	-0.1295 ***	0.0012	-0.0870 ***	0.0012	-0.0231 ***	0.0010
7	-0.1089 ***	0.0012	-0.0697 ***	0.0012	-0.0130 ***	0.0011
8	-0.0919 ***	0.0012	-0.0554 ***	0.0012	-0.0048 ***	0.0011
9	-0.0814 ***	0.0013	-0.0474 ***	0.0013	-0.0015	0.0011
10	-0.0704 ***	0.0013	-0.0389 ***	0.0013	0.0030 **	0.0012
11	-0.0629 ***	0.0014	-0.0343 ***	0.0013	0.0023	0.0012
12	-0.0494 ***	0.0014	-0.0237 ***	0.0014	0.0074 ***	0.0012
13	-0.0409 ***	0.0014	-0.0185 ***	0.0014	0.0089 ***	0.0013
14	-0.0338 ***	0.0015	-0.0131 ***	0.0015	0.0105 ***	0.0013
15	-0.0284 ***	0.0015	-0.0096 ***	0.0015	0.0113 ***	0.0013

	Base size	Average size	Firm age only	Base size with firm age controls	Average size with firm age controls
Diagnostic statistics					
Number of observations (millions)	10.5	10.5	10.5	10.5	10.5
R-squared	0.1795	0.0188	0.4855	0.4936	0.5847

* significantly different from reference group at the 0.05 level.

** significantly different from reference group at the 0.001 level.

*** significantly different from reference group at the 0.0001 level.

1. Reference group is 1,000 or more, except for "Firm age only."

2. Reference group is 16 or older, except for "Base size" and "Average size."

Note: ILU = Individual Labour Units. Standard errors are correct for heteroscedasticity. Year and two-digit North American Industry Classification System (NAICS) industry controls are also included in these regressions.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Table 7
Regression results, using individual labour units — Continuing firms

Variables	Base size		Average size		Firm age only		Base size with firm age controls		Average size with firm age controls	
	estimate	standard error	estimate	standard error	estimate	standard error	estimate	standard error	estimate	standard error
Firm size¹ (ILU)										
< 1	0.4071 ***	0.0029	-0.0629 ***	0.0028	0.3361 ***	0.0029	-0.1600 ***	0.0027
1 to <5	0.0235 ***	0.0028	-0.0126 ***	0.0026	0.0049	0.0028	-0.0594 ***	0.0026
5 to < 10	-0.0262 ***	0.0028	-0.0009	0.0026	-0.0300 ***	0.0028	-0.0262 ***	0.0026
10 to < 20	-0.0233 ***	0.0028	0.0049	0.0026	-0.0243 ***	0.0028	-0.0131 ***	0.0026
20 to < 50	-0.0147 ***	0.0028	0.0092 ***	0.0026	-0.0145 ***	0.0028	-0.0032	0.0026
50 to < 100	-0.0061 *	0.0029	0.0119 ***	0.0027	-0.0054	0.0029	0.0033	0.0027
100 to < 250	-0.0034	0.0030	0.0121 ***	0.0028	-0.0028	0.0030	0.0061 *	0.0027
250 to < 500	-0.0023	0.0034	0.0101 **	0.0032	-0.0024	0.0034	0.0049	0.0032
500 to < 1,000	-0.0009	0.0041	0.0042	0.0039	-0.0016	0.0041	-0.0001	0.0038
Firm age² (years)										
1	0.3076 ***	0.0008	0.2239 ***	0.0007	0.3410 ***	0.0008
2	0.0099 ***	0.0007	-0.0168 ***	0.0007	0.0307 ***	0.0007
3	0.0026 ***	0.0007	-0.0199 ***	0.0007	0.0212 ***	0.0007
4	0.0032 ***	0.0007	-0.0167 ***	0.0007	0.0199 ***	0.0007
5	0.0000	0.0007	-0.0176 ***	0.0007	0.0152 ***	0.0007
6	-0.0001	0.0007	-0.0161 ***	0.0007	0.0140 ***	0.0007
7	0.0003	0.0007	-0.0143 ***	0.0007	0.0134 ***	0.0007
8	0.0004	0.0007	-0.0133 ***	0.0007	0.0128 ***	0.0007
9	0.0008	0.0008	-0.0119 ***	0.0008	0.0123 ***	0.0008
10	0.0018 *	0.0008	-0.0097 ***	0.0008	0.0125 ***	0.0008
11	0.0004	0.0008	-0.0100 ***	0.0008	0.0101 ***	0.0008
12	0.0017 *	0.0008	-0.0075 ***	0.0008	0.0104 ***	0.0008
13	0.0003	0.0009	-0.0073 ***	0.0009	0.0081 ***	0.0009
14	0.0000	0.0009	-0.0072 ***	0.0009	0.0072 ***	0.0009
15	0.0005	0.0009	-0.0060 ***	0.0009	0.0071 ***	0.0009

	Base size	Average size	Firm age only	Base size with firm age controls	Average size with firm age controls
Diagnostic statistics					
Number of observations (millions)	8.1	8.1	8.1	8.1	8.1
R-squared	0.0606	0.0035	0.0451	0.0846	0.0515

* significantly different from reference group at the 0.05 level.

** significantly different from reference group at the 0.001 level.

*** significantly different from reference group at the 0.0001 level.

1. Reference group is 1,000 or more, except for "Firm age only."

2. Reference group is 16 or older, except for "Base size" and "Average size."

Note: ILU = Individual Labour Units. Standard errors are correct for heteroscedasticity. Year and two-digit North American Industry Classification System (NAICS) industry controls are also included in these regressions.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.


The parameters on the dummies for each firm size class are plotted in Charts 4 to 7 for both the base-year and average-year sizing methods, including and not including age dummies. Three results stand out. The first is that, regardless of method, mean employment growth rates at the higher end of the size distribution differ very little.

The second result is that a negative relationship between firm size and employment growth prevails only for the base-year method and only for the smallest firms. Employment growth rates of small firms are greater only when regression to the mean is ignored. And even then, the relationship disappears at relatively small firm sizes (5 to 10 employees). The average-year method yields a positive relationship over the same interval; that is, small firms have no advantage. The difference between the results of the two methods suggests that regression-to-the-mean matters, at least for the smallest firms.³¹

The third result is that, as in HJM, firm age matters. With the base-year method, including age reduces or eliminates the negative relationship between firm size and employment growth rates, depending on whether all firms or only continuing firms are considered. The higher growth rates of the smallest firms are attributable to the younger age of this group. With the average-year method, the positive relationship holds over a wider range of firm sizes when firm age is controlled, but disappears beyond the 10-to-20-employee range. Here, growth rates actually rise with increasing firm size, at least for the smaller size classes. In both cases, smaller firms do not have proportionately higher growth rates.

Chart 4

Relationship between growth and firm size, all firms, 1999 to 2008 — Average Labour Units




Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

31. The regressions were also performed using Birch's methodology—classifying firms according to their base year size, with the growth rate defined as $(X_{t-1} - X_{t-2})/X_{t-2}$ instead of $(X_t - X_{t-1})/\text{avg}(X_t, X_{t-1})$. Only continuing firms were considered, because the standard growth rate is not well defined for entries and exits. None of the coefficients on size classes are significant whether or not age is included, because the standard errors are too large.

Chart 5


Relationship between growth and firm size, all firms, 1999 to 2008 — Individual Labour Units



Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Chart 6


Relationship between growth and firm size, continuing firms, 1999 to 2008 — Average Labour Units



Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Chart 7

Relationship between growth and firm size, continuing firms, 1999 to 2008 — Individual Labour Units




Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

The parameters on the age variables are presented in Charts 8 to 11. The vast majority of estimates are significant at 0.01%, indicating a difference between the employment growth rates of the omitted group (age 16 or older) and younger firms. Since growth rates for births are 2 by construction, they are excluded from the chart. Two features stand out. First, controlling for firm size does not influence the nature of the relationship between firm age and employment growth, although it does influence its magnitude. In all cases, employment growth rates decline sharply between the first and second year of operation.³² When both continuing and exiting firms are included, a pronounced positive relationship between firm age and employment growth after the second year is evident. However, for continuing firms, there is only a small negative relationship after the second year. This indicates a higher exit rate for young firms compared with old firms. HJM call this the "up or out" dynamics of start-ups and very young firms.

32. HJM also find a decrease in growth rates between the first and second year, though of a smaller magnitude. The pronounced drop observed here is partly due to the part-year births that survive into year one. Their growth rate in year one might be overestimated.

Chart 8


The relationship between growth and age, all firms — Average Labour Units



Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Chart 9


The relationship between growth and age, all firms — Individual Labour Units



Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Chart 10


**The relationship between growth and age, continuous firms —
Average Labour Units**



Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Chart 11

**The relationship between growth and age, continuous firms —
Individual Labour Units**



Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Together, the results suggest that short-run regression-to-the-mean effects matter in Canada. Little evidence supports a relationship between firm size and employment growth, except for the very smallest firms. On the contrary, the results suggest that Gibrat's Law holds for mean employment growth rates of firms at or above minimum efficient scale; that relationship sees higher rates of employment growth as firm size increases toward minimum efficient scale; and what relationship does exist is driven by births and very young firms, with little difference among older firms.

These results hold whether employment is measured by ALU or ILU. This suggests that firms modify quality-adjusted and quality-unadjusted employment in a similar way year-on-year. Idiosyncratic shocks, therefore, have a similar effect on these two concepts of employment.

4.3 Size distribution

Additional evidence supporting the applicability of Gibrat's Law to the employment growth rates of Canadian businesses is provided by the size distribution of Canadian firms. If Gibrat's Law holds, firm size should follow a Zipf distribution, with a Zipf exponent of -1.³³ Whether the Zipf exponent ξ is significantly different from -1 was tested by regressing the log of the rank of the firm (ordered from largest to smallest) on its log size (Gabaix, 1999):

$$\log(\text{rank}) = \log a + \xi \log(\text{size})$$


If Zipf's Law holds, firms should fall along a straight line with a slope of -1.

Chart 12 suggests that Zipf's Law is a good approximation for firms with 5 employees or more (log of firm size of 0.7). Smaller firms, however, clearly do not fit the pattern.

33. "This necessary emergence of Zipf's law may sound surprising. An analogy for it would be the central limit theorem: if a variable of arbitrary distribution (of finite variance) is examined and the mean of its successive realizations is calculated, normalized appropriately, this mean will always have (asymptotically) a normal distribution, independently of the characteristic of the initial process. Likewise, whatever the particulars driving the growth . . . , as soon as they satisfy (at least over a certain range) Gibrat's law, [the] distribution will converge to Zipf" (Gabaix, 1999).

Chart 12**Zipf regression results illustrated, using Average Labour Units**

Log of firm rank



Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Table 8 presents the Zipf regression coefficients for firms active in 2008. At the aggregate level, when all firms are included, the coefficient is significantly different from -1. Restricted to firms with 5 or more employees, the coefficient is only slightly above 1 using both ALU (-1.03) and ILU (-1.06) measures of employment. These results are close to the coefficient of -1.06 found by Axtel for the United States (Axtel, 2001). Thus, Zipf's Law is a good approximation of the upper tail of the distribution of firm size for Canada.

Table 8**Zipf regression results, 2008**

	Intercept		Zipf exponent		Observations number
	estimate	standard error	estimate	standard error	
All firms					
Average Labour Units	5.71 ***	0.00017	-0.56 ***	0.00023	1,046,300
Individual Labour Units	5.89 ***	0.00018	-0.68 ***	0.00024	1,046,300
Firms with 5 employees or more					
Average Labour Units	6.18 ***	0.00023	-1.03 ***	0.00018	252,200
Individual Labour Units	6.28 ***	0.00017	-1.06 ***	0.00014	315,200

*** significantly different from 0 at the 0.0001 level.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.



5 Conclusion

Which firms grow and which shrink is an increasingly important question. In countries with high unemployment, especially in the wake of the financial crisis, policy-makers benefit from knowing what kinds of firms create jobs.

An extensive literature has examined the relationship between firm size and employment growth. Many studies find a negative relationship. Others, however, argue that the evidence in support of this result relies on inadequate data, omitted variables (firm age), and misleading interpretations, and report little or no relationship between the number of employees in a firm and its propensity to create jobs.

This paper analyzed the employment growth patterns of Canadian firms using Statistics Canada's LEAP database. The results show that regression-to-the-mean effects are important, at least for the smallest firms. Age also matters, especially for the youngest firms. When the influence of transitory shocks is taken into account and when firm age is controlled, little evidence emerges to support the contention that small firms have proportionately higher employment growth rates. Gibrat's Law holds for the mean growth rate; that is, above a 20-employee threshold, little relationship is evident between firm size and employment growth. These results are qualitatively similar to those of Haltiwanger et al. (2010).

The size distribution of firms that results from the pattern of employment growth rates was also examined. Gibrat's Law should produce a distribution of firms into size categories that follows Zipf's Law—a small number of large firms and a large number of small firms. And in fact, the size distribution of Canadian firms is well approximated by a Zipf distribution, which is consistent with the independence of growth rates and firm size.

It is sometimes claimed that small firms grow faster than large firms. Others argue that this is not so. This paper finds that both positions are correct. It is true that average growth rates are larger for a broad swath of firms defined arbitrarily as smaller. But it is not true that as we move across the spectrum of size classes from the very smallest to the very largest that growth rates decline monotonically with size. Indeed, there is evidence that average growth rates increase initially. Having proportionately more smallish firms in this circumstance will not lead to proportionately more job growth.

Even this finding needs to be qualified because growth rate averages conceal considerable differences across the underlying population brought about by industry composition and year effects. Once these are considered using multivariate analysis, there is little in the way of a relationship between job growth and size of the firm, especially after age is considered. That still means that taken together small firms create proportionately more jobs but that much of this occurs in the population studied because particular industries where small firms are heavily represented were growing more quickly, or the time period studied was one with a particular concentration of years where small firms did better than large firms, or because there was a particular concentration of younger firms that were small and it is age that is closely related to growth. What was driving growth in the aggregate of small firms were industry-specific factors,

time-specific circumstances or the age of the participants. The circumstances encouraging growth are to be found in these areas, not the size of a firm per se.

This study also points to a new set of research questions on the start-up process and early development. For instance, what are the characteristics of start-ups and young firms that are successful? Do they respond to opportunities differently than their failing competitors? Related to this is the topic of firms smaller than minimum efficient scale (MES). A positive relationship between firm size and growth rates was found for firms below 20 employees. This positive relationship is still found even when only continuing firms are considered. Being closer to this size not only improves the survival probability of firms below it but also is associated with higher employment growth. Further studies could be devoted to examining how growth is affected by the initial distance of entrants from this size and how it relates to MES in general.

Another topic that deserves attention is the variance of growth rates. The literature for Canada and other countries finds that the variance of growth rates is not independent of firm size. Therefore, Gibrat's Law might not hold for variances in Canada even though it holds for means. This issue could be studied with the database developed in this paper, which would also allow the exploration of the relationship between the variance of growth rates and firm age.

6 Appendix

Figure 1
Merger example

	Employment in	
	t-1	t
Small firm (A)	100	0
<i>Firm B purchases Firm A</i>		
Large firm (B)	600	725
<hr/>		
<u>Job growth without accounting for mergers and acquisitions</u>		
Small firm (A)		-100
Large firm (B)		125
Net growth		<hr/> 25
 <u>Firm growth in LEAP</u>		
Large firm (B)		25
Net growth		<hr/> 25
<i>Firm A no longer appears in the data</i>		

Note: When a small firm is acquired by a large firm, the Longitudinal Employment Analysis Program (LEAP) does not record a large decrease in employment in the small firm and a large increase in the large firm. Instead, the acquiring firm is assigned the annual job growth of the now merged entity, and the acquired firm is no longer present in the data.

6.1 Comparison between Average Labour Unit (ALU) and Individual Labour Unit (ILU)

Table 9

Average individual labour unit to average labour unit ratio, by average labour unit firm size


Firm size (average labour units)	Observations	Average ILU/ALU ratio
	number	ratio
Less than 1	3,494,100	5.7
1 to less than 5	3,648,200	1.6
5 to less than 10	1,028,600	1.3
10 to less than 20	637,700	1.2
20 to less than 50	457,100	1.1
50 to less than 100	151,700	1.1
100 to less than 250	77,400	1.1
250 to less than 500	18,100	1.1
500 to less than 1,000	8,600	1.1
1,000 or more	8,700	1.1
All firms	9,530,100	3.0

Note: Based on cross-sectional observations with positive ALU and ILU in final year.

Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Chart 13


Total average labour units and individual labour units in business sector, 1999 to 2008




Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.

Chart 14

Employment dynamics for the business sector using average labour units and individual labour units, 1999 to 2008



Source: Statistics Canada, Longitudinal Employment Analysis Program, 1999 to 2008.



References

- Audretsch, D.B., L. Klomp, E. Santarelli and A.R. Thurik. 2004. "Gibrat's Law: Are the services different?" *Review of Industrial Organization*. Vol. 24. No. 3. p. 301–324.
- Axtell, R.L. 2001. "Zipf distribution of U.S. firm sizes." *Science*. Vol. 293. No. 5536. p. 1818–1820.
- Baldwin, J.R. 1998. "Were small producers the engines of growth in the Canadian manufacturing sector in the 1980s?" *Small Business Economics*. Vol. 10. No. 4. p. 349–364.
- Baldwin, J.R., R. Dupuy and W. Penner. 1992. "Development of longitudinal panel data from business registers." The Canadian experience. *Statistical Journal of the United Nations*. Vol. 9. p. 289–303.
- Baldwin, J.R., and W. Gu. 2011. "Firm dynamics and productivity growth: A comparison of the retail trade and manufacturing sectors." *Industrial and Corporate Change*. Vol. 20. No. 2. p. 367–395.
- Baldwin, J.R., and G. Picot. 1994. *Employment Generation by Small Producers in the Canadian Manufacturing Sector*. Statistics Canada Catalogue no. 11F0019M. Ottawa, Ontario. Analytical Studies Branch Research Paper Series. No. 70.
- Baldwin, J.R., and A. Lafrance. 2011. *Firm Turnover and Productivity Growth in Selected Canadian Services Industries, 2000 to 2007*. Statistics Canada Catalogue no. 11F0027M. Ottawa, Ontario. Economic Analysis (EA) Research Paper Series. No. 72.
- Baldwin, J.R., and J. Haltiwanger. 1998. "A comparison of job creation and job destruction in Canada and the United States." *Review of Economics and Statistics*. Vol. 80. No. 3. p. 347–356.
- Bartelsman, E., S. Scarpetta and F. Schivardi. 2003. *Comparing Analysis of Firm Demographics and Survival: Micro-level Evidence for the OECD Countries*. Paris. OECD Publishing. OECD Economics Department Working Papers. No. 348. <http://dx.doi.org/10.1787/010021066480> (accessed on February 21, 2012).
- Birch, D. L. 1981. "Who creates jobs?" *The Public Interest*. Vol. 65. p. 3–14.
- Davis, S.J., J. Haltiwanger and S. Schuh. 1993. *Small Business and Job Creation: Dissecting the Myth and Reassessing the Facts*. Cambridge, Massachusetts. National Bureau of Economic Research (NBER). NBER Working Papers. No. 4492.
- Drolet, M., and R. Morissette. 1998. *Recent Canadian Evidence on Job Quality by Firm Size*. Statistics Canada Catalogue no. 11F0019M. Ottawa, Ontario. Analytical Studies Branch Research Paper Series. No. 128.
- Eurostat/Organization for Economic Cooperation and Development (OECD). 2007. *Eurostat-OECD Manual on Business Demography Statistics*. Paris. OECD Publications. <http://www.oecd.org/dataoecd/8/8/39974460.pdf> (accessed on February 21, 2012).

- Evans, D.S. 1987. "The relationship between firm growth, size, and age: Estimates for 100 manufacturing industries." *Journal of Industrial Economics*. Vol. 35. No. 4. p. 567–581.
- Fujiwara, Y., C.D. Guilmi, H. Aoyama, M. Gallegati and W. Souma. 2003. "Do Pareto-Zipf and Gibrat Laws hold true? An analysis with European firms." *Physica A: Statistical Mechanics and its Applications*. Vol. 335. No. 1-2. p. 197–216.
- Gabaix, X. 1999. "Zipf's Law for cities: An explanation." *The Quarterly Journal of Economics*. Vol. 114. No. 3. p. 739–767.
- Gabaix, X. 2009. "Power laws in economics and finance." *Annual Review of Economics*. Vol. 1. No. 1. p. 255–294.
- Geroski, P.A. 2005. "Understanding the implications of empirical work on corporate growth rates." *Managerial and Decision Economics*. Vol. 26. No. 2. p. 129–138.
- Haltiwanger, J.C., R.S. Jarmin and J. Miranda. 2010. *Who Creates Jobs? Small vs. Large vs. Young*. Cambridge, Massachusetts. National Bureau of Economic Research (NBER). NBER Working Papers. No. 16300.
- Lafrance, A., and D. Leung. 2010. *T2-LEAP (1984–2007): A Longitudinal Database of Incorporated Firms in Canada*. Ottawa, Ontario. Analytical Studies Branch, Statistics Canada. Technical report. Unpublished.
- Luttmer, E.G.J. 2007. "Selection, growth, and the size distribution of firms." *The Quarterly Journal of Economics*. Vol. 122. No. 3. p. 1103–1144.
- Mata, J. 2008. "Gibrat's Law". *The New Palgrave Dictionary of Economics*. Second edition. S.N. Durlauf and L. Blume (eds.). Vol. 3. p. 662–664.
- Neumark, D., B. Wall and J. Zhang. 2008. *Do Small Businesses Create More Jobs? New Evidence for the United States from the National Establishment Time Series*. Cambridge, Massachusetts. National Bureau of Economic Research (NBER). NBER Working Papers. No. 13818.
- Okolie, C. 2004. "Why size class methodology matters in analyses of net and gross job flows." *Monthly Labor Review*. Vol. 127. No. 7. p. 3–12.
<http://www.bls.gov/opub/mlr/2004/07/art1full.pdf> (accessed on February 21, 2012).
- Petrunia, R. 2008. "Does Gibrat's Law hold? Evidence from Canadian retail and manufacturing firms." *Small Business Economics*. Vol. 30. No. 2. p. 201–214.
- Statistics Canada. 2011. *Employment, Earnings and Hours*. Statistics Canada Catalogue no. 72-002-X. Vol. 89. No. 11. Ottawa, Ontario.
- Sutton, J. 1997. "Gibrat's legacy." *Journal of Economic Literature*. Vol. 35. No. 1. p. 40–59.
- Tornqvist, L., P. Vartia and Y. Vartia. 1985. "How should relative change be measured?" *The American Statistician*. Vol. 39. No. 1. p. 43–46.